Journal of Archaeological Science 53 (2015) 156–165

Contents lists available at ScienceDirect

Journal of Archaeological Science

journal homepage: http://www.elsevier.com/locate/jas

Observed methods of cuneiform tablet reconstruction in virtual and real world environments

Andrew Lewis [a,*], Sandra Woolley [a], Eugene Ch'ng [b,c], Erlend Gehlken [d]

[a] University of Birmingham, UK
[b] School of Computer Science, International Doctoral Innovation Centre, University of Nottingham, Ningbo, China
[c] Centre for Creative Content and Digital Innovation, University of Malaya, Malaysia
[d] Universität Frankfurt/Main, Germany

A R T I C L E I N F O

Article history:
Received 6 May 2014
Received in revised form
26 September 2014
Accepted 30 September 2014
Available online 8 October 2014

Keywords:
Collaboration
3D visualization
Virtual environments
Fragment reassembly
Artefact reconstruction
Cuneiform

A B S T R A C T

The reconstruction of fragmented artefacts is a tedious process that consumes many valuable work hours of scholars' time. We believe that such work can be made more efficient via new techniques in interactive virtual environments. The purpose of this research is to explore approaches to the reconstruction of cuneiform tablets in the real and virtual environment, and to address the potential barriers to virtual reconstruction of fragments. In this paper we present the results of an experiment exploring the reconstruction strategies employed by individual users working with tablet fragments in real and virtual environments. Our findings have identified physical factors that users find important to the reconstruction process and further explored the subjective usefulness of stereoscopic 3D in the reconstruction process. Our results, presented as dynamic graphs of interaction, compare the precise order of movement and rotation interactions, and the frequency of interaction achieved by successful and unsuccessful participants with some surprising insights. We present evidence that certain interaction styles and behaviours characterise success in the reconstruction process.

1. Introduction

There are a considerable number of cuneiform tablets and fragments in the collections of the world's museums. Most of the tablets originate from Mesopotamia, the land between the rivers Tigris and Euphrates which cover modern day Iraq, parts of Syria and Turkey. The cuneiform tablets were formed of clay taken from the river banks. The cuneiform script is characterized by wedge shaped impressions on the surface of the clay tablets due to the form of the reed stylus which was used to write the texts. Cuneiform tablets vary in both width and length. A survey of tablets (Lewis and Ch'ng 2012) in the Cuneiform Digital Library Initiative database (CDLI) showed that most tablets ranged from 20 to 60 mm in size, although some tablets are larger.

As would be expected from cultures at the height of their development, the cuneiform texts convey a wide range of information, including religious texts, literature, mathematics, astronomy, medicine, law, letters, royal decrees, contemporary events, educational matters, and administrative documents like inventories and orders, bills, contracts as well as certificates of authenticity from traders. The intellectual diversity of the tablet contents is matched by the variation of the tablet size and condition. This paper explores issues specific to the field of physical and virtual cuneiform reconstruction, and suggests a system capable of assisting with the reconstruction of cuneiform tablets using virtual representations of cuneiform fragments.

Projects like the Cuneiform Digital Library Initiative (http://cdli.ucla.edu), the Cuneiform Digital Forensic Project (CDFP) (Woolley et al. 2002), and the BDTNS (Database of Neo-Sumerian Texts - http://bdts.filol.csic.es/) have advanced the process of cataloguing cuneiform collections in the digital realm, and brought collected resources of museums and universities onto the desktop computer. This has resulted in a reduction in the time required to search cuneiform archives for text. A networked computer can search through thousands of text fragments in a fraction of a second, and draw results from multiple resources regardless of geographical location.

Unfortunately, the process of cuneiform tablet reconstruction has not been affected so positively by the advancement of

* Corresponding author. Digital Humanities Hub, University of Birmingham, Edgbaston, Birmingham, B15 2TT, United Kingdom.
E-mail addresses: axl148@bham.ac.uk, andrew@monkeysailor.co.uk (A. Lewis).

http://dx.doi.org/10.1016/j.jas.2014.09.024
0305-4403/

technology, and the processes employed to rebuild broken cuneiform tablets still rely on glue and putty. Manual joining of fragments from catalogue descriptions and pieces in individual collections are still the prevalent methods of reconstruction. This is partly because existing digital databases pay particular attention to the textual content of a fragment rather than its exact physical dimensions, which can make reuniting broken fragments very difficult for individuals without specific training or access to the original fragments. More importantly, there are limited tools available that allow for the digital capture and intuitive manipulation of scanned 3D fragments in a virtual environment.

The virtual reconstruction of cuneiform fragments presents a two-fold problem. Firstly, the fragments presented on screen must be sufficiently well defined for a user to examine in detail and make decisions about placement. The shape of the individual fragments must be easy to identify when viewed on screen in proximity to other similar fragments, and the surface of the fragments should be of a sufficient resolution to allow close examination from multiple viewpoints. Secondly, the nature of the reconstruction task requires fine manipulation of fragments, and a suitable interface for this task must be considered. As Poupyrev et al. (1997) explain, the manipulation of objects in virtual environments can be awkward and inconvenient because of the lack of tactile feedback and other interface considerations.

With respect to the problems of representation and reproduction, scholars working with cuneiform texts have relied until now on manual observation and interpretation of the physical evidence at hand. Whilst these scholars have been diligent in their task, there has always existed the possibility for error and misinterpretation.

In the case of purely lithographic representations of cuneiform tablets, the chances of transcription and substitution errors have existed throughout the publishing pipeline, as was noted by the past Keeper of Egyptian and Assyrian Antiquities in the British Museum, E. A. Wallis Budge (1925). Even photographic representations cannot guarantee a robust representation of fragments, because the camera orientation, position, and lighting can all affect the clarity and apparent geometry of the object (Hameeuw and Willems, 2011). The advent of high-resolution flatbed scanners and digital photography has led to the digitization of cuneiform fragments and the foundation of international online databases like the CDLI and the Database of Neo-Sumerian Texts BDTNS. Unfortunately, the principal issue of legibility when representing a 3D shape in a 2D medium remains unsolved. The problem of accurate representation has been discussed for well over 100 years, and one article in The Journal of the Photographic Society of London in 1866 gave specific reference to the difficulties of representing cuneiform text (Diamond, 1864).

Research has demonstrated the potential of the technology for 3D cuneiform representation (Woolley et al. 2001; Willems et al., 2005), and Anderson and Levoy (2002) suggested the use of 3D visualization and scanning techniques in the analysis of complete cuneiform tablets. Anderson and Levoy also provide useful technical information about minimum resolution requirements for the accurate reproduction of cuneiform tablets with legible text, and although the paper deals primarily with tablets that have already been reconstructed, the arguments in favour of 3D representation are still valid for cuneiform fragments. Cohen et al. (2004) and Hahn et al. (2007) made use of 3D scanning and visualization technology in the digital Hammurabi project, which produced high resolution textured scans of tablets, while Levoy's advocacy of 3D scanning and visualization techniques continued in the 2006 paper “Fragments of the City: Stanford's Digital Forma Urbis Romae Project”. In this paper, Levoy explains how fragments of the Forma Urbis Romae (an 18 m long map of Rome produced circa 206 CE) were laser scanned and reconstructed using inscribed surface topology and fragment edges. Their paper also discusses the value of manual tagging of topographic features as a key for future reconstructions.

There is evidence that 3D scanning can provide appropriate virtual representations and open the field of virtual reconstruction to the automated techniques of computer assisted reconstruction seen with skull fragments in the fields of bioarchaeology, palaeoanthropology, and skeletal biology (Gunz et al. 2009; Kuzminsky and Gardiner, 2012), and also with pot and plasterwork in the fields of pot and fresco reconstruction (Brown et al. 2010; Karasik and Smilansky, 2008; Laugerotte and Warzée, 2004; Papaioannou et al. 2002). The wider academic community provides many examples where an increased understanding of a subject has resulted from the analysis of 3D data. The in situ analysis of engravings in archaeological sites (Güth, 2012), the analysis and reconstruction of coins and coin fragments in numismatics (Zambanini et al. 2009, 2008), and the capture of graffiti on Roman pottery (Montani et al., 2012) are representative cases. More generally, the application of techniques for the automatic recording and illustration of artifacts (Gilboa et al. 2013) could be applied to 3D cuneiform models, and used to streamline the process of documentation while removing one potential source of recording error. More specific techniques for the reconstruction of cuneiform tablets have been made in Ch'ng et al. 2013 and Lewis and Ch'ng 2012, which include the analysis of the complete tablet size as a template for fragment reconstruction, and the use of stigmergy as a model for interaction between users.

Furthermore, it is possible that many generalized algorithms could be adapted to select or orient particular fragments for reconstruction (Kleber and Sablatnig, 2009; Demaine and Demaine, 2007). For example, the popularity of Optical Character Recognition (OCR) software has ensured that a number of language independent methods exist for recognizing the orientation of written data (Hochberg et al. 1995; Lu and Tan, 2006), and it is probable that these can be adapted to suit the cuneiform text found on the tablets. Analysis of the fractal dimension (Wong et al. 2005) of an edge might also provide a useful index for sorting potentially matching edges.

The capture and visualization of fragments represents only one part of the virtual cuneiform reconstruction problem. Manipulation of fragments in virtual space is an issue that must be considered, and it is likely that initial tests with a virtual environment will give mixed results when users with variable experience engage with a 3D interface for the first time. Keehner (2006) and Vora et al. (2002) indicate that participation in virtual tasks has a positive learning effect, and dexterity will improve as interaction continues. Other issues, such as the lack of depth perception and haptic feedback are less easy to address. 3D visualization presents one possible avenue for investigation, as for example, stereo 3D has been shown to increase attention and offer a more natural interactive experience (Schild et al. 2012), but caution must be exercised because increased visual fatigue and even nausea may occur after prolonged use (Yu et al., 2012). Newer gestural interfaces like the LeapMotion™ or Microsoft Kinect™ may also be considered as novel methods for interaction, but at this time they lack sufficient resolution for stable manipulation of fragments. Electromechanical polymer screens (Kim et al. 2013) and holographic haptic devices (Iwamoto et al. 2008) may in the future be able to provide tactile surface feedback to users. The detail of the matching surfaces of an artefact are usually so complex that anything less than a high resolution physical reproduction of the fragments such as those produced, for example, by the Creative Machines laboratory at Cornell University (Knapp et al. 2008) would be of limited value in the haptic sense.

The advances in related fields such as fresco reconstruction and pottery reconstruction suggest that the problems caused by virtual

abstraction are not insurmountable, but in order to overcome them we must first investigate the interaction issues specific to cuneiform fragment reassembly.

2. Materials and methods

With the exception of Ch'ng et al. (2013) which suggests that a solution to the problems associated with cuneiform reconstruction may exist in the field of complexity science, there is currently no published research specific to cuneiform reconstruction strategy. The first goal of the research presented here was to determine some of the basic techniques employed by participants to match together and to discard clay fragments in both the real and virtual world. To achieve this, five sets of clay tablet fragments were scanned using a NextEngine HD 3D scanner. Each set contained between 6 and 8 fragments which were scanned in at medium resolution (at 2.5 k sample points per inch), with each model containing approximately 1.5 million vertices. The resulting models were decimated to reduce the vertex count to approximately 30 thousand vertices and were then imported into a custom made virtual 3D environment (Vizard based) configured to accept mouse and keyboard input to control the position and rotation of the fragments in virtual space. The application also supported stereoscopic 3D visualization using an interlaced field pattern and polarized glasses. A computer with an AMD Phenom II x4 955 processor, 8 Gb of RAM, and an Nvidia GTX 560i graphics card was used for each test. A generic 105 key QWERTY keyboard and a 3 button optical mouse with scroll wheel were connected as input devices, and an LG Cinema 3D Monitor (D2342P) was used for both 2D and 3D output.

Pilot studies were carried out to determine appropriate time limits for reconstruction tasks in the virtual and physical environments during each experiment. From these pilot studies it was determined that a time limit of 12 min was appropriate for virtual tasks. After consideration from multiple sources (Bertaux, 1981; Guest et al. 2006; Mason, 2010; Marshall, 1996; Neilsen & Landauer 1993; Schmettow, 2012), it was decided that as the current study represented a precursor to a larger investigation and involved both qualitative and quantitative aspects, sufficient information to determine the direction of future work could be obtained with a relatively small number of participants. In total, 15 participants performed the experiments, 8 of which were male and 7 were female. The mean age of participants was 32 years, with the youngest participant being aged 24 and the oldest age was 41. Each participant was isolated for the duration of the test in the Chowen Prototyping Hall at the University of Birmingham, and presented with a series of tasks involving three methods of interaction (see Illustration 1):

1. Physical reconstruction task

   The participant was asked to reconstruct physical tablets from a collection or collections of fragments. Participants were informed at the beginning of each task that the collection of fragments they were presented with may be pieces from one tablet, more than one tablet, or may not fit together at all. The collections were sorted so that they contained the fragments of a complete tablet and either zero or more superfluous fragments. The purpose of this task was to provide baseline values for current reconstruction methods, and explore the effect of superfluous fragments on the manual reconstruction process.

2. Virtual reconstruction task

   Participants were presented with the equivalent reconstruction tasks of physical participants, but were given virtual 3D fragments rather than their real-world counterparts.

3. Stereoscopic virtual reconstruction task

   Participants were shown virtual fragments on a 3D monitor, and asked to perform the same reconstruction tasks as described above. This test restores a sense of depth perception to the participant, but still requires manipulation of 3D objects using standard input devices. This separates the effects of the lost depth perception from the effects of remote object manipulation using a keyboard and mouse.

Participants were also asked to reconstruct sets that contained either 2 superfluous fragments, or a number of superfluous fragments equal to the number of valid fragments (N) in the set. These tasks were referred to as $N+2$ and $2N$ respectively. In all cases, the time taken to complete the task and the accuracy of the completed tablet were recorded, as was the time to make the 1st and 2nd join. For virtual tasks, the physical operations (rotate, move) used to achieve the end result were recorded in a log of participant interactions during each test. At the completion of each task, the participant was asked a series of questions to elicit qualitative feedback. The environment used in the experiments was consistent, with physical surfaces coloured black to match the background colour of the screen used in the virtual tasks. Identical input and output devices were used for all virtual tasks, and instructions were provided in a script. Information about the controls for the

Illustration 1. Screenshot showing virtual reconstruction task on the left, in contrast to a physical reconstruction task on the right.

virtual system were provided on a printed sheet next to the computer, which the participant was instructed to read before the test began. The sheet remained in place next to the computer for the duration of the experiment.

3. Experimental results

All participants in the first test group were able to reconstruct the physical fragments into complete tablets well within the allotted time. The fastest join (*i.e*. the time to join the first two fragments together) was made within 5 s with the average time to the first join being 34.6 s. The average time between the first and second match was 33.8 s. The fastest participant completed the entire process within 65 s. No participant took more than 5 min and 49 s to reconstruct the tablet from the set of fragments that they were given.

The interaction methods employed by participants fell into two broad categories: *Methodical* and *Selective*. *Methodical* interactions involved a "brute-force" approach to the reconstruction process, comparing fragments systematically and then retaining those pieces that join together. *Selective* interactions were more discriminating, involving careful observation of the fragments before choosing those that were likely to form a cogent pair. It was observed that participants favoured a particular method of interaction, and did not tend to change their method. It was also observed that the manual manipulation of fragments was very free, with multiple simultaneous operations. It was not unusual for rotation and movement operations to be carried out in both hands at the same time. The initial freedom of motion became compromised as the number of fragments being held increased, so that participants were forced to discard the collections that they were holding in order to manipulate only relevant pieces. This became problematic as the reconstructed tablets neared completion. Several participants commented that glue or tape would have been helpful during the reconstruction process. Contrarily, the deliberate exclusion of simulated gravity from the virtual environment means that holding fragments in position is not an issue, although some participants noted that a method of grouping individual fragments into a single object would have made manipulation easier. Unfortunately, the restrictions of a virtual interface using standard equipment currently prevent the fluid ambidextrous manipulation of multiple fragments. When using a keyboard and mouse, the participant is restricted to sequential actions on a single fragment, which in turn increases the time required to manipulate fragments into the desired position.

Performance in the virtual tasks was significantly lower than in the physical, with only one of the participants managing to reconstruct a complete tablet before the end of the 12 min session. However, 11 of the 15 of participants were able to make at least one successful join, with the fastest participant taking 27 s to make a connection. Another participant had the shortest inter-match time (the time between a participant making the first and second join), taking just 33 s to find the second join (see Figs. 1 and 2).

With the sequential nature of virtual manipulation (where users are restricted by the interface into performing actions on only one fragment at a time), almost 75% of the actions carried out by the participant are rotations, which typically occur before a participant moves fragments together.

The participant interactions were classified so that participants who were able to make at least two matches in the virtual system were deemed to be *successful*, while those who made fewer than two joins were classed as *unsuccessful*. Successful participants typically rotated fragments less, with an average of approximately 72%, ranging between 56% and 83% rotations. In contrast, 77% of the interactions made by unsuccessful participants were rotations, ranging between 70% and 92%.

Fig. 3 shows the rotation and translation events for a particular participant over the course of the experiment. The numerical identifier of the fragment being manipulated is expressed on the *Y* axis, with the time in seconds progressing along the *X* axis. The participant's actions shown in Fig. 3 illustrate a heavy bias towards fragment rotation. These participants were unable to find any matches between the fragments, and ultimately stopped without making a single match. In comparison, Fig. 4 shows the activity of more successful participants who made at least two joins from the provided set. These participants manoeuvred the fragments into

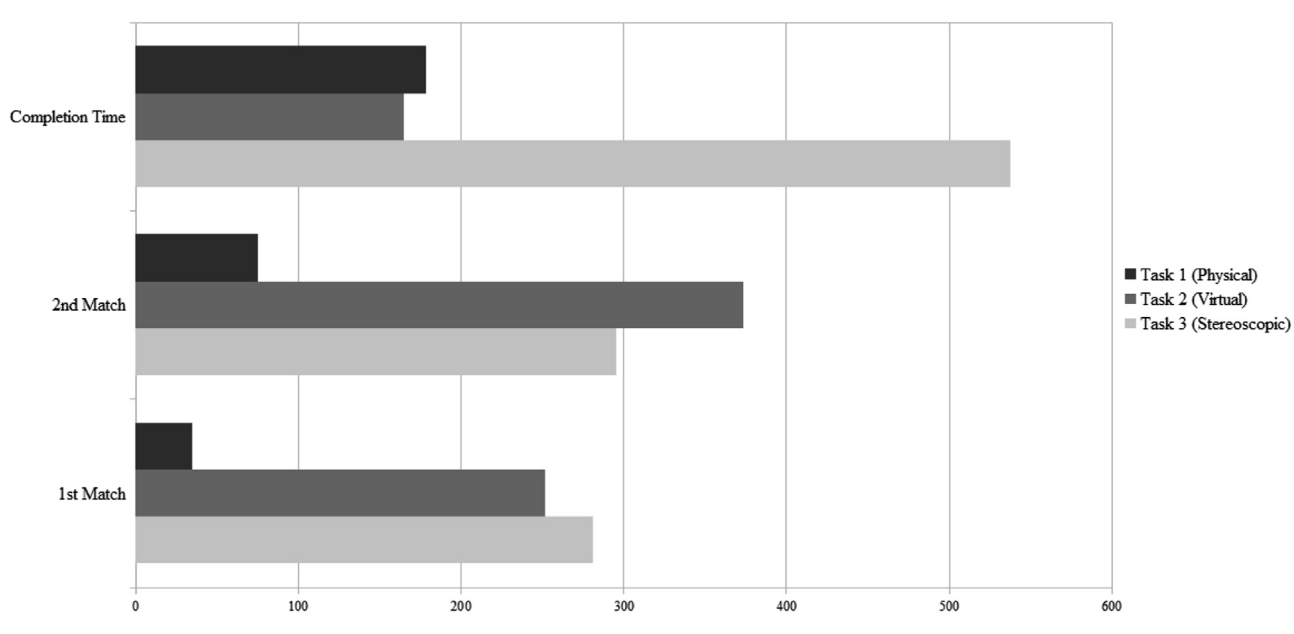

Fig. 1. Graph showing the mean 1st match, 2nd match and completion time for each task.

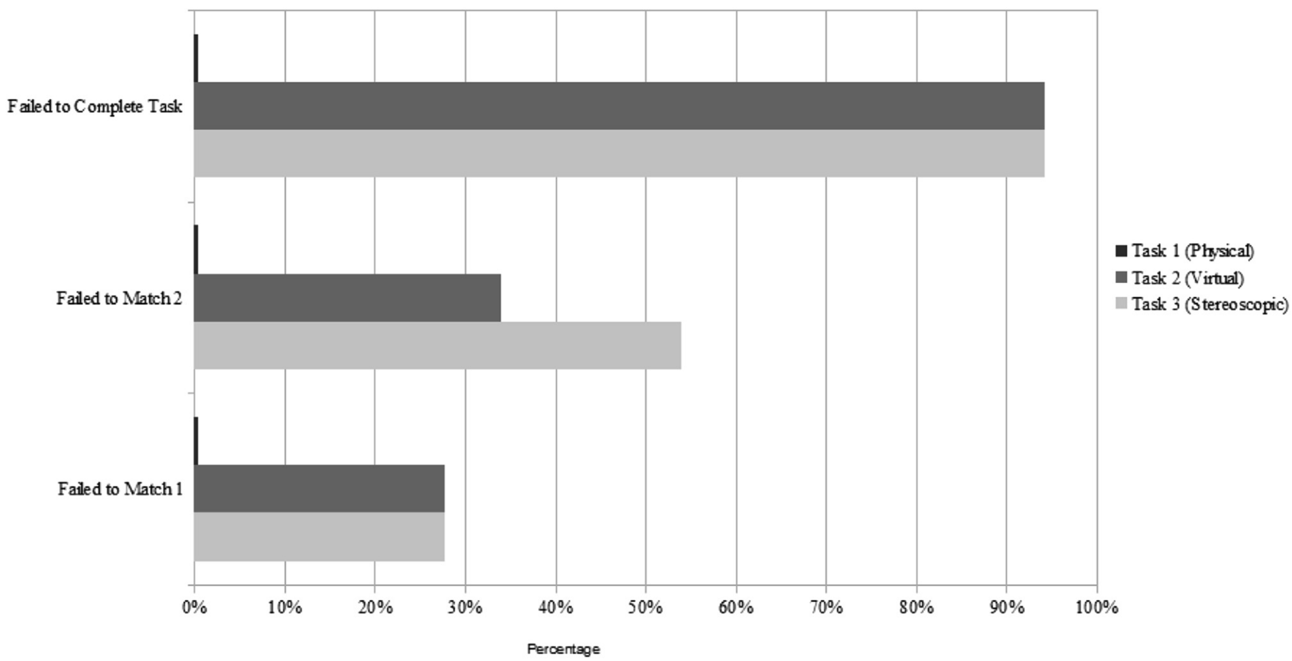

Fig. 2. Graph showing percentage of participants unable to reach experimental milestones for each task.

close proximity after an initial inspection, and then continued to manipulate them until they were either matched or discarded.

If a participant aligns one fragment so that the edge appears to join with another fragment, the participant will move the fragments together and attempt a close fit. Pieces that do not match will typically be moved away from the target piece and discarded. This method of virtual reconstruction is reminiscent of the selective strategy employed by some participants in the manual reconstruction experiments. It is possible that the speed reduction encountered when using the virtual interface makes a brute-force,

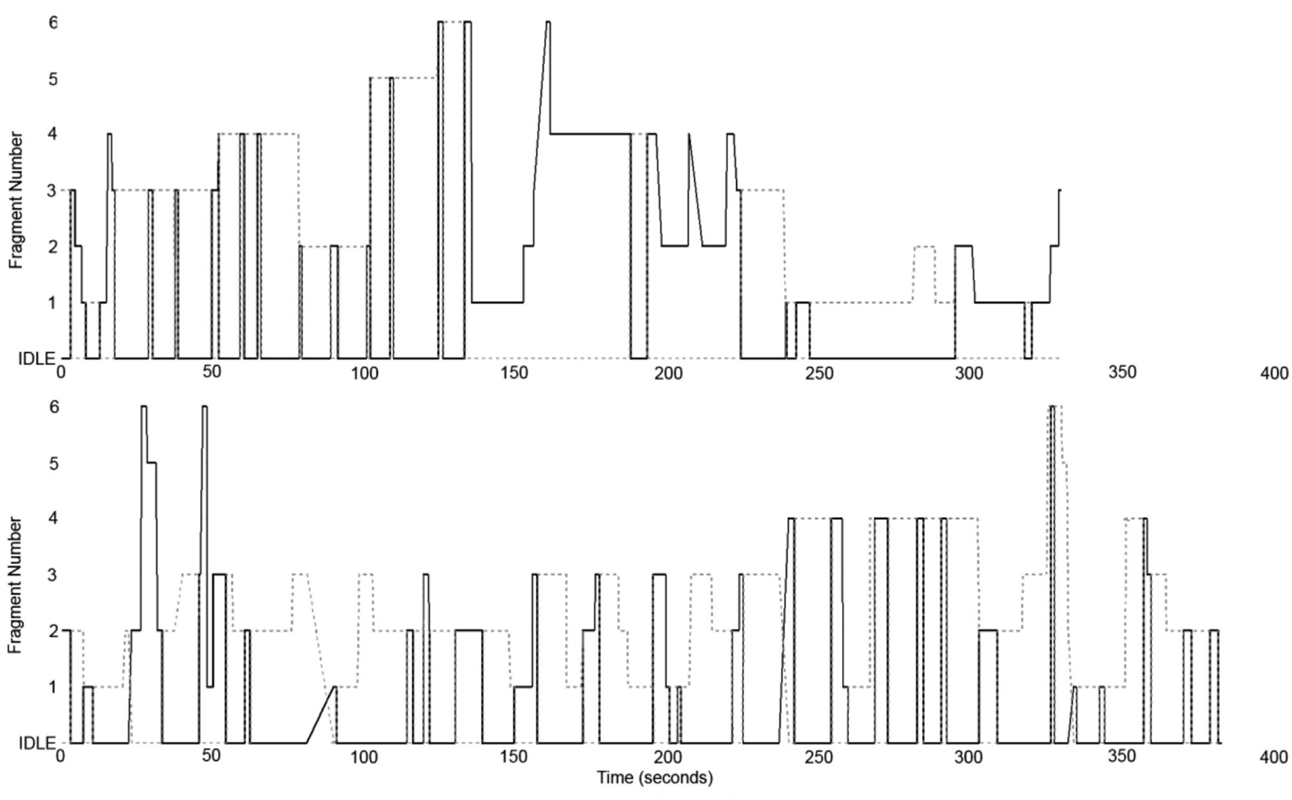

Fig. 3. Graph showing the rotation and movement actions of unsuccessful participants when using the virtual reconstruction system.

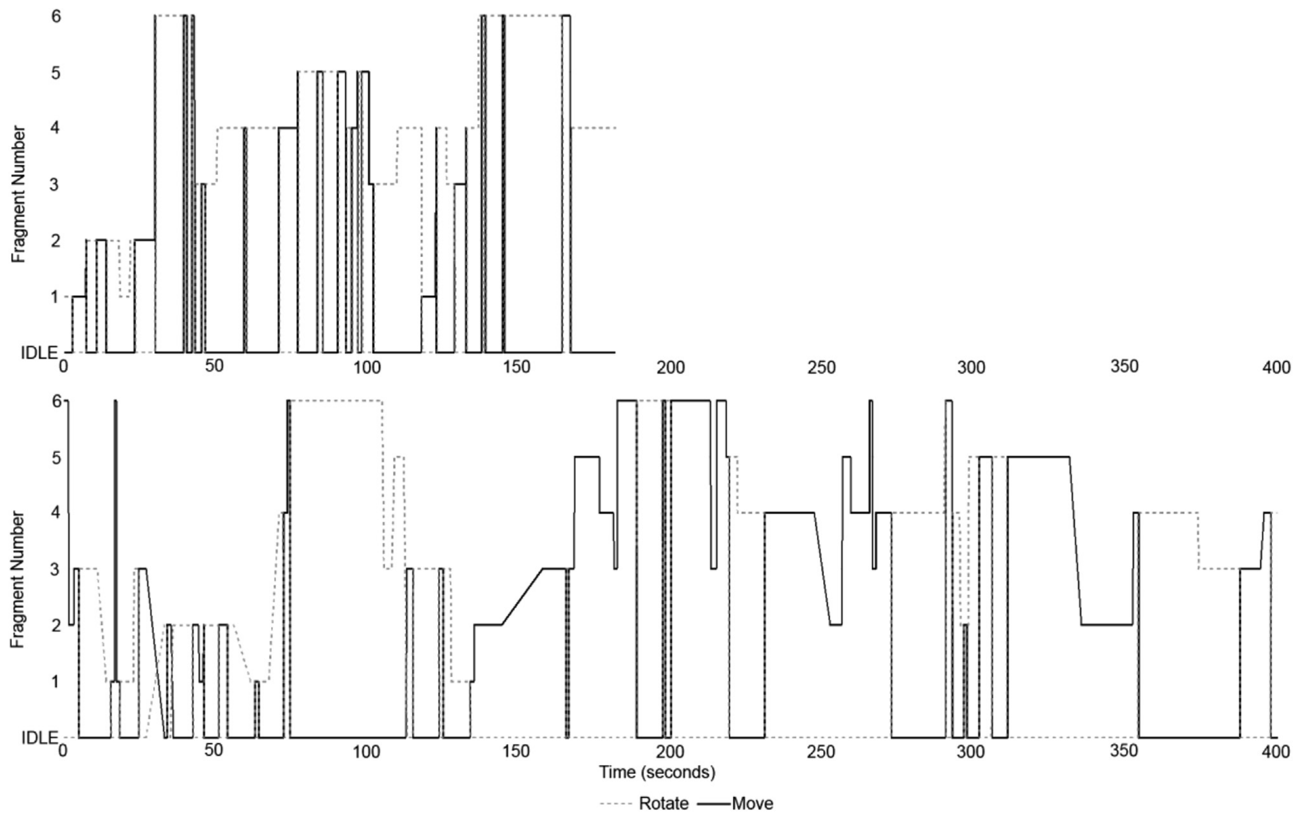

Fig. 4. Graph showing the rotation and movement actions of successful participants when using the virtual reconstruction system.

methodical approach to the joining process too laborious for users to focus on.

In common with physical strategy, 14 of the 15 participants began their digital reconstruction tasks by manipulating one of the larger fragments in the set, with 6 participants choosing the largest available fragment regardless of its position on screen. This mirrors observational evidence from the physical tests and also the feedback from several users on their individual reconstruction strategies.

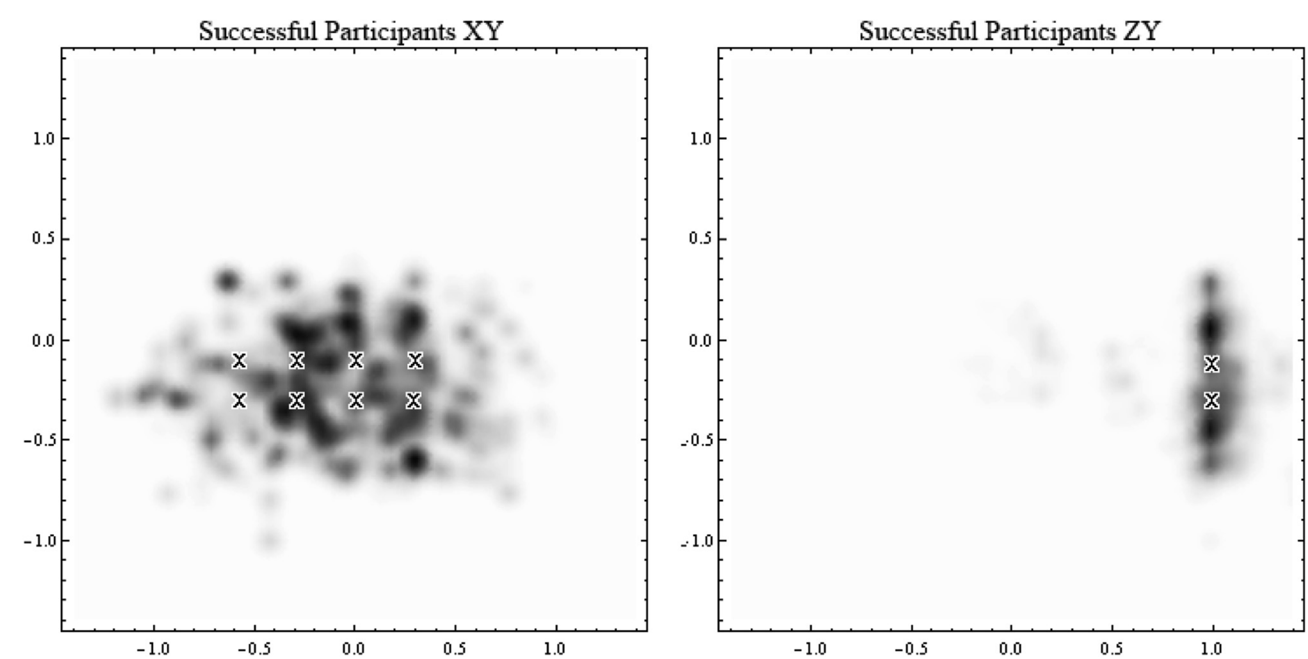

Fig. 5. Interaction map showing the average frequency of fragment interaction in 3D space for successful participants. The left hand graph represents a "screen view", whilst the right hand graph shows the depth of fragments within the space. Crosses indicate the starting position of fragments.

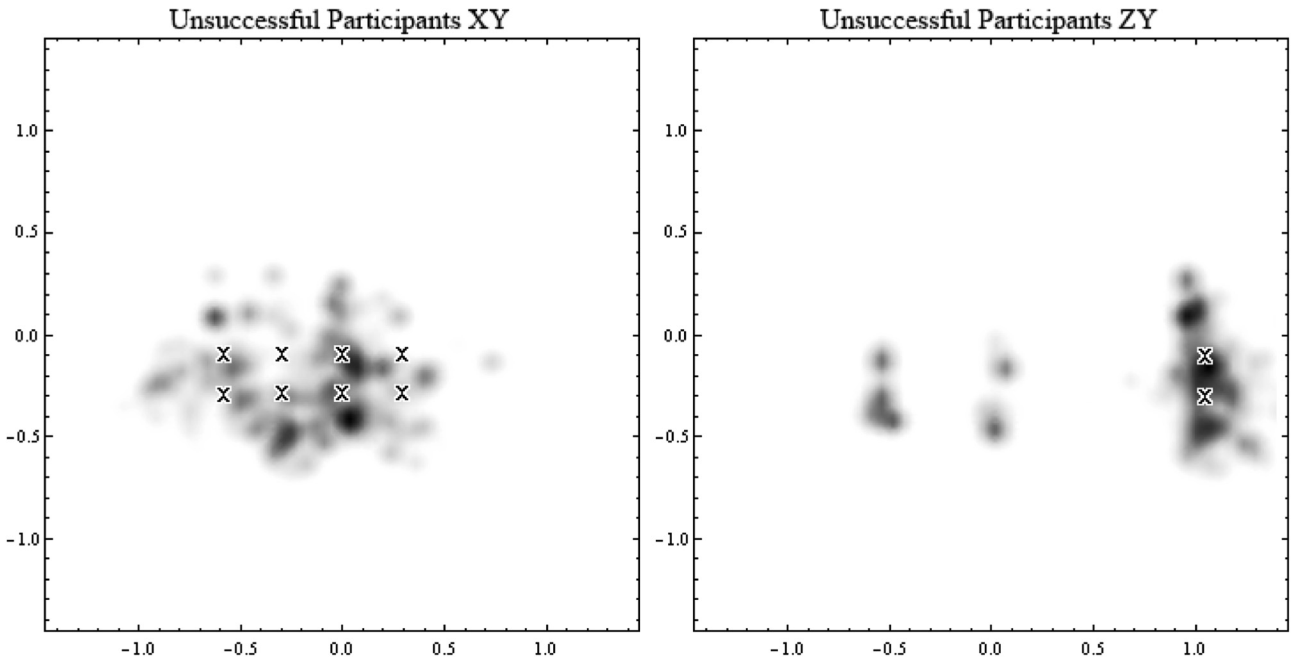

Fig. 6. Interaction map showing the average frequency of fragment interaction in 3D space for unsuccessful participants. The left hand graph represents a "screen view", whilst the right hand graph shows the depth of fragments within the space. Crosses indicate the starting position of fragments.

The size of the first fragment chosen by the user did not directly affect the speed at which the participants made matches, although it may be useful to consider this preference for starting when designing a virtual system that can automatically suggest fragments to users. In the majority of these cases, the users will be looking for a smaller fragment than the one they currently hold.

Graphing the points of interaction within the virtual space reveals that unsuccessful participants (those who made fewer than two joins in the virtual system) were more likely to pull fragments towards the camera to enlarge them, while successful participants (those who made two or more joins in the virtual system) spent more time interacting with fragments at their original location.

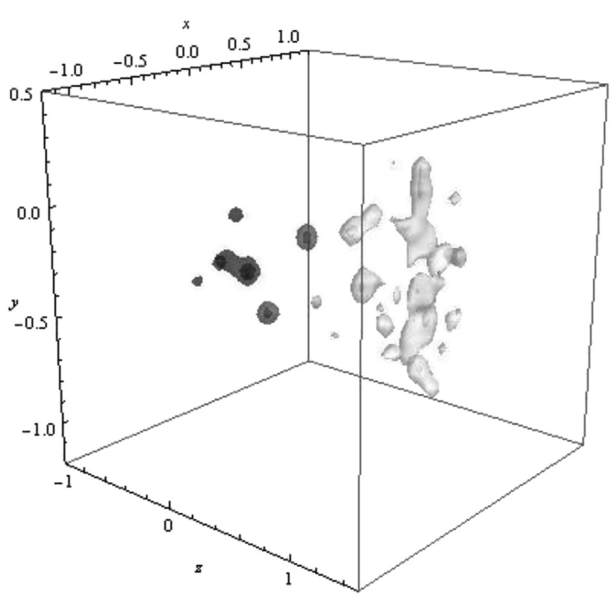

Fig. 7. Graph showing the interaction patterns of unsuccessful participants in the virtual space.

These interaction maps in Figs. 5 and 6 show a front (*XY*) and side (*ZY*) view of the virtual space, with the areas of most activity being shaded darker. If we examine these graphs, we can see that the most noticeable clusters of activity are at depth 1 in the *Z* axis, which is the default starting position that fragments are placed on the screen.

This activity is present for both successful and unsuccessful participants. The graph of the unsuccessful participants also shows clusters of activity at depth 0 and at −0.5 which indicates that the fragments have been moved towards the camera. The disparity between the interactions of the successful and unsuccessful participants is more pronounced when viewed in 3D.

Fig. 7 is a 3D representation of this spatial interaction information and shows the sparse interaction patterns of the unsuccessful participants, with isolated areas of activity towards the default fragment depth of 1 and the zero point of the graph. In contrast, the successful participants whose activities are illustrated in Fig. 8 show a greater level of activity at the default fragment depth, whilst very little activity occurs in other areas of the virtual space.

As would be expected, the introduction of superfluous fragments appears to increase the time that participants need to make a match, with the minimum completion time increasing as the number of spurious fragments increases. This is reflected in the results from the physical tasks as shown in Fig. 9.

4. Discussion

Participants revealed several key features that could be used to improve the virtual reconstruction process. Recurrent attributes identified by participants include the surface markings and colour of a fragment. The smoothness of fragment surface was also identified as allowing participants to distinguish sign areas and blank surface areas from obviously broken edges. Participants commented that the size of the fragments was important, with larger fragments being used as anchor points for testing smaller fragments against. This was also shown in the analysis of the logs of initial interaction with fragment sizes from the virtual

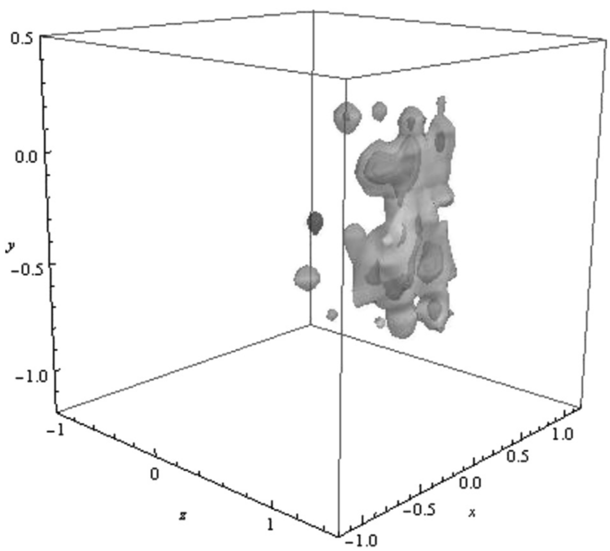

Fig. 8. Graph showing the interaction patterns of successful participants in the virtual space.

environment. Virtually pre-sorting larger collections of fragments by these features may improve efficiency of reconstruction. This technique has seen some success in the field of fresco reconstruction, and a virtual system to suggest fragments based on these features is the next logical step.

Many subjects stated that the lack of haptic (tactile) feedback was an issue during the virtual reconstruction process, and the lack of depth perception (leading to problems with object scaling) was also mentioned by multiple users. While the effect of depth perception was investigated during this study, the effect of haptic feedback and touch were less easy to test at this stage. A larger study has been planned to investigate the effectiveness of touch screen technology and explore several alternative techniques for interaction and visualization on static and mobile platforms.

It was assumed that the early performance of the participants in the virtual tasks would depend in part on their previous exposure to 3D software, and those participants with previous experience of 3D modelling and GIS software would be more comfortable manipulating objects in 3D space from the beginning. This proved not to be the case, which tallies with the results of other experiments and suggests that a longer exposure to the virtual interface over a course of multiple sessions would improve the performance of participants in the reconstruction tasks.

The 3D heatmaps reveal that the interactions of successful participants in perpendicular planes (i. e. in our experiments in planes parallel to the *XY*plane, see Fig. 7) occur over a wider area than those of unsuccessful participants, while motion at different points on the *Z*-axis is less frequent. The interactions of unsuccessful participants exhibit a greater range of motion along the *Z* axis, with less overall motion in planes parallel to the *X*−*Y* plane. We see from this that successful participants make more use of the available *X*−*Y* screen space, with more activity occurring in the spaces between hotspots. In contrast, the unsuccessful participants have a much less energetic profile, with more separation in the *Z* axis. It is possible that the effect of perspective scaling is a contributing factor in the performance of these participants, with distant fragments being misinterpreted as smaller than they actually are.

Multiple participants commented that virtual reconstruction was more difficult because the depth of the fragments was indeterminate, and pieces that appeared to fit together were actually positioned at different depths, although this was not apparent on the 2D screen. While the use of binocular 3D subjectively increased the effectiveness of the virtual reconstruction environment, it produced no measurable positive effect to the reconstruction process, and had negative associations with the availability of the technology and the increased eye fatigue caused by convergence/fixed-focus. One participant was unable to work with the 3D screen despite having no binocular vision defects. Several participants claimed to feel more able to perform the task when working with stereoscopic 3D models, but ultimately performed no better than those working with normal screens. In measured terms, fewer

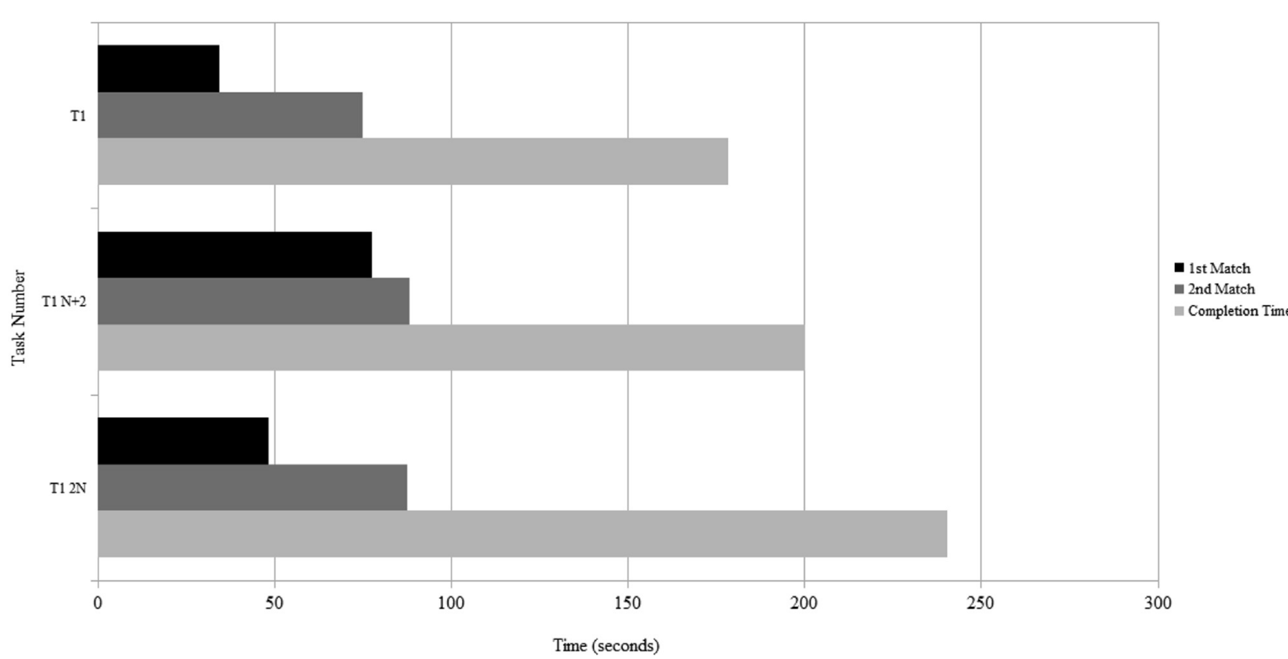

Fig. 9. The effect of additional fragments on reconstruction time for participants in task 1.

participants were able to make a second join when using stereoscopic 3D within the allotted time, but overall their performance was on par with participants working without stereoscopic glasses.

Participants also stated that the lack of tactile feedback was a significant drawback for virtual reconstruction. While it may currently be impossible to implement accurate tactile feedback within the virtual system, it is possible that additive manufacturing techniques could be used to provide a physical copy of fragments that appear to join in the virtual system. These printed fragments could then be used to make a definitive decision on the validity of a proposed join. More extensive use of additive printing technology could also be considered so that staff with limited training can carry out multiple fitting operations concurrently. Replica parts are low value and replaceable, having no special handling requirements or storage considerations.

5. Concluding remarks

In the course of our experiments, we observed several behaviours that could improve the virtual reconstruction process for cuneiform fragments. Firstly, we observed that more successful participants kept fragments close to each other in the *Z* axis, and as such a visual representation of *Z* depth within the workspace may help to help participants to perform better. However, we observed that restoring depth perception by stereographic representation does not improve participant performance. We have also observed that participants tend to begin with a larger fragment, with which they then try to match with smaller fragments. In a virtual system that automatically suggests possible matches, a bias toward suggesting smaller fragments than the one currently held may also improve the participant's performance. The absence of tactile feedback was noted by several users, and while no technology currently exists to completely restore the sense of tactility, it may be possible to provide an audio or visual feedback system that provides feedback on the closeness of fit between multiple fragments. One example of such a system might be a border around the visible fragment that becomes more opaque as the closeness of fit between the fragments increases. Other features that could improve the experience for participants working within a virtual system include the ability to glue multiple fragments together so that they can be manipulated as a single object, and the ability to magnify fragments so that close inspection of edges can be carried out quickly. The results of our experiments indicate that the manual reconstruction of fragments is faster than virtual reconstruction, but the physical world does not allow for easy parallel processing of fragment sets, nor does it permit casual accessibility. Despite the limitations of a virtual system, the potential for task parallelization and human computation makes virtual reconstruction an attractive choice for fragment joining.

Crowdsourcing projects like the Galaxy Zoo (http://www.galaxyzoo.org/) which use human volunteers to classify new images of galaxies, and Cellslider (https://www.zooniverse.org/project/cellslider) which uses a similar framework to identify potentially cancerous cells, provide a platform for the classification of scientific images that computers are currently unable to match. These projects show how crowdsourcing can be used successfully for human computation, with existing tools like CrowdCurio and Zooniverse being able to connect potential participants with researchers for free (http://www.zooniverse.org, http://www.crowdcurio.com). Other services like Amazon Mechanical Turk (http://www.mturk.com/mturk/) provide a framework for participants to bid and work on a variety of projects in exchange for money. The success of these projects suggests another potential method for the reconstruction of artefacts, with a virtual environment providing an interface for paid or voluntary human workers. If the ethical considerations of wages estimated in the range of US$ 1.25 per hour for Mechanical Turk (Ross et al. 2010), the lack of worker's rights (Fort et al. 2011), and potential security concerns can be avoided, the potential power of crowdsourcing is difficult to dismiss.

A distributed system designed to maximize the advantages of the virtual environment whilst minimizing the inherent limitations could open up the field of cuneiform reconstruction to new audiences, and free scholars from the drudgery of manual reconstruction. It is also likely that the research behind such a system would be applicable to a number of other fields within the archaeological community.

Acknowledgements

This project gratefully acknowledges the support of The Leverhulme Trust (project grant number F000 94 BP), the Chowen Prototyping Hall, and the multidisciplinary support of the Heritage and Cultural Learning Hub at the University of Birmingham.

References

Amazon Mechanical Turk, (accessed 01.02.14.). http://www.mturk.com/mturk.
Anderson, S., Levoy, M., 2002. Unwrapping and visualizing cuneiform tablets. IEEE Comput. Graph. Appl. 22, 82–88.
Bertaux, D., 1981. Biography and Society: the Life History Approaches in the Social Sciences. Sage Publications Ltd.
Brown, B.J., et al., 2010. Tools for virtual reassembly of fresco fragments. In: Seventh International Conference on Science and Technology in Archaeology and Conservations.
Budge, E., 1925. The Rise and Progress of Assyriology. M. Hopkinson & Co. Ltd.
Cellslider, (accessed 01.02.14.). https://www.zooniverse.org/project/cellslider.
Ch'ng, E., Lewis, A.W., Ghelken, E., Woolley, S., 2013. A theoretical framework for stigmergetic reconstruction of ancient text. In: Ch'ng, E., Gaffney, V.L., Chapman, H. (Eds.), Visual Heritage in the Digital Age, Springer Cultural Computing Series.
Cohen, J., Duncan, D., Snyder, D., Cooper, J., Kumar, S., Hahn, D., Chen, Y., Purnomo, B., Graettinger, J., 2004. iClay : digitizing cuneiform. In: The 5th International Symposium on Virtual Reality, Archaeology and Cultural Heritage (VAST 2004), Brussels, pp. 135–143.
CrowdCurio, (accessed 01.02.14.). http://www.crowdcurio.com.
Cuneiform Digital Library Initiative, (accessed 01.02.14.). http://cdli.ucla.edu/.
Database of Neo-Sumerian Texts, (accessed 01.02.14.). http://bdts.filol.csic.es/.
Demaine, E., Demaine, M., 2007. Jigsaw Puzzles, edge matching, and Polyomino Packing: connections and complexity. Graphs Comb. 23 (1), 195–208.
Diamond, H.W., 1864. Photography: importance of the art. In: The Journal of the Photographic Society of London, No. 143, pp. 139–140. Available at: http://books.google.co.uk/books/about/THE_JOURNAL_OF_THE_PHOTOGRAPHIC_SOCIETY.html?id=IDEFAAAAQAAJ.
Fort, K., Adda, G., Bretonnel, C.K., 2011. Amazon mechanical turk: gold mine or coal mine? Comput. Linguist. 37 (2), 413–420.
Galaxy Zoo, (accessed 01.02.14.). http://www.galaxyzoo.org/.
Gilboa, A., Tal, A., Shimshoni, I., Kolomenkin, M., 2013. Computer-based, automatic recording and illustration of complex archaeological artifacts. J. Archaeol. Sci. 40 (2), 1329–1339.
Guest, G., Bunce, A., Johnson, L., 2006. How many interviews are enough?: an experiment with data saturation and variability. Field Methods 18, 59–82.
Gunz, P., Mitteroecker, P., Neubauer, S., Weber, G.W., Bookstein, F.L., 2009. Principles for the virtual reconstruction of hominin crania. J. Hum. Evol. 57 (1), 48–42.
Güth, A., 2012. Using 3D scanning in the investigation of Upper Palaeolithic engravings: first results of a pilot study. J. Archaeol. Sci. 39, 3105–3114.
Hahn, D., Baldwin, K., Duncan, D., 2007. Non-laser-based scanner for three-dimensional digitization of historical artifacts. Appl. Opt. 46 (15), 2838–2850.
Hameeuw, H., Willems, G., 2011. New visualization techniques for cuneiform texts and sealings. Akkadica 132 (2), 163–178.
Hochberg, J., Kerns, L., Kelly, P., Thomas, T., 1995. Automatic script identification from images using cluster-based templates. IEEE Trans. Pattern Anal.Mach. Intell. 19, 176–181.
Iwamoto, T., Tatezono, M., Shinoda, H., 2008. Non-contact method for producing tactile sensation using airborne ultrasound. In: Proc. EuroHaptics 2008, pp. 504–513.
Karasik, A., Smilansky, U., 2008. 3D scanning technology as a standard archaeological tool for pottery analysis: practice and theory. J. Archaeol. Sci. 35 (5), 1148–1168.
Keehner, M., 2006. Learning a spatial skill for surgery: how the contributions of abilities change with practice. Appl. Cogn. Psychol. 20 (4), 487–503.

Kim, S.-C., Israr, A., Poupyrev, I., 2013. Tactile rendering of 3D features on touch surfaces. In: Proc. Of UIST, vol. 13, pp. 23–30.
Kleber, F., Sablatnig, R., 2009. A survey of techniques for document and archaeology artefact reconstruction. In: Proceedings of International Conference on Document Analysis and Recognition, vol. 10, pp. 1061–1065.
Knapp, M., Wolff, R., Lipson, H., 2008. Developing printable content: a repository for printable teaching models. In: Proceedings of the 19th Annual Solid Freeform Fabrication Symposium, Austin TX.
Kuzminsky, S.C., Gardiner, M.S., 2012. Three dimensional laser scanning: potential uses for museum conservation and scientific research. J. Archaeol. Sci. 39, 2744–2751.
Laugerotte, W., Warzée, N., 2004. An environment for the analysis and reconstruction of archaeological objects. In: VAST 2004: the 5th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage, pp. 165–174.
Lewis, A., Ch'ng, E., 2012. A photogrammetric analysis of cuneiform tablets for the purpose of digital reconstruction. Int. J. Cult. Herit. Digit. Era, EuroMED Suppl. 1 (1), 49–53.
Lu, S., Tan, C., 2006. Automatic document orientation detection and categorization through document vectorization. In: Proceedings of the 14th ACM International Conference on Multimedia, pp. 113–116.
Marshall, M., 1996. Sampling for qualitative research. Fam. Pract. 13, 522–526.
Mason, M., 2010. Sample size and saturation in Phd studies using qualitative interviews. Forum: Qual. Soc. Res. 11 (3). Article 8.
Montani, I., Sapin, E., Sylvestre, R., Marquis, R., 2012. Analysis of roman pottery graffiti by high resolution capture and 3D laser profilometry. J. Archaeol. Sci. 39, 3349–3353.
Nielsen, J., Landauer, T., 1993. A mathematical model of the finding of usability problems. In: Proceedings of ACM INTERCHI'93 Conference, pp. 206–213.
Papaioannou, G., Karabassi, E., Theoharis, T., 2002. Reconstruction of three-dimensional objects through matching of their parts. IEEE Trans. Pattern Anal. Mach. Intell. 24, 114–124.
Poupyrev, I., Weghorst, S., Billinghurst, M., Ichikawa, T., 1997. A framework and testbed for studying manipulation techniques for immersive VR. In: Virtual Reality Software and Technology. Proceedings of the ACM Symposium on Virtual Reality Software and Technology, pp. 21–28.
Ross, J., Irani, L., Silberman, M., Zaldivar, A., Tomlinson, B., 2010. Who are the crowdworkers?: shifting demographics in mechanical turk. In: CHI '10 Extended Abstracts on Human Factors in Computing Systems (CHI EA '10), pp. 2863–2872.
Schild, J., LaViola, J., Masuch, M., 2012. Understanding user experience in stereoscopic 3D games. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12), pp. 89–98.
Schmettow, M., 2012. Sample Size in Usability Studies in Communications of the ACM, vol. 55(4), pp. 64–70.
Vora, J., Nair, S., Gramopadhye, A.K., Duchowski, A.T., Melloy, B.J., Kanki, B., 2002. Using virtual reality technology for aircraft visual inspection training: presence and comparison studies. Appl. Ergon. 33 (6), 559–570.
Willems, G., Verbiest, F., Moreau, W., Hameeuw, H., Van Lerberghe, K., Van Gool, L., 2005. Easy and cost-effective cuneiform digitizing. In: Proceedings of the 6th International Symposium on Virtual Reality, Archaeology and Cultural Heritage, pp. 73–80.
Wong, A., Wu, L., Gibbons, P., 2005. Fast estimation of fractal dimension and correlation integral on stream data. Inf. process. Lett. 91–97.
Woolley, S.I., Flowers, N.J., Arvanitis, T.N., Livingstone, A., Davis, T.R., Ellison, J., 2001. 3D capture, representation, and manipulation of cuneiform tablets. In: Corner, B.D., Nurre, J.H., Pargas, R.P. (Eds.), Three-dimensional Image Capture and Applications IV, Proceedings of SPIE, vol. 4298, pp. 103–110.
Woolley, S.I., Davis, T.R., Flowers, N.J., Pinilla-Dutoit, J., Livingstone, A., Arvanitis, T.N., 2002. Communicating cuneiform: the evolution of a multimedia cuneiform database. J. Visible Lang. Spec. Ed. Res. Commun. Des. 36 (3), 308–324.
Yu, J.-H., Lee, B.-H., Kim, D.-H., 2012. EOG based eye movement measure of visual fatigue caused by 2D and 3D displays. In: IEEE-EMBS International Conference on Biomedical and Health Informatics (BHI), pp. 305–308.
Zambanini, S., Schlapke, M., Kampel, M., Müller, A., 2008. On the use of computer vision for numismatic research. In: The 9th International Symposium on Virtual Reality, Archaeology and Cultural Heritage VAST (2008).
Zambanini, S., Schlapke, M., Kampel, M., Müller, A., 2009. Historical coins in 3D: acquisition and numismatic applications. In: The 10th International Symposium on Virtual Reality, Archaeology and Cultural Heritage (VAST 2009).
Zooniverse, (accessed 01.02.14.). http://www.zooniverse.org.